



CIRRELT

Centre interuniversitaire de recherche
sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

The Vehicle Routing Problem with Pauses

Jean-Philippe Gagliardi
Jacques Renaud
Angel Ruiz
Leandro C. Coelho

May 2014

CIRRELT-2014-22

Document de travail également publié par la Faculté des sciences de l'administration de l'Université Laval,
sous le numéro FSA-2014-05.

Bureaux de Montréal :

Université de Montréal
Pavillon André-Aisenstadt
C.P. 6128, succursale Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :

Université Laval
Pavillon Palais-Prince
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirreлт.ca

The Vehicle Routing Problem with Pauses

Jean-Philippe Gagliardi, Jacques Renaud, Angel Ruiz, Leandro C. Coelho*

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)
and Department of Operations and Decision Systems, 2325 de la Terrasse, Université Laval,
Québec, Canada G1V 0A6

Abstract. In this paper we introduce the Vehicle Routing Problem with Pauses (VRPP). This problem arises when drivers must make pauses during their shift, for example, for lunch breaks. Drivers breaks have already been considered in long haul transportation when drivers must rest during their travel, but their technical and scientific aspects have been mostly neglected in the less than truckload and last mile distribution contexts up to date. In this paper we introduce the first mathematical formulation for the VRPP. This linear formulation has the disadvantage of roughly doubling the number of nodes, and thus significantly increasing the size of the distance matrix and the number of variables. Consequently, standard branch-and-bound algorithms are only capable of solving very small-sized instances. In order to tackle large instances, we propose a fast local search-based heuristic tailored for the VRPP, which is shown to be very efficient. Through a series of computational experiments, we show that solving the VRPP without explicitly considering the pauses during the optimization can lead to a number of infeasibilities and to higher solution costs. These results demonstrate the importance of integrating drivers pauses in the resolution process.

Keywords. Vehicle routing problem with pauses, exact formulation, heuristic.

Acknowledgements. This research was partially supported by Grants OPG 0293307 and OPG 0172633 from the Natural Sciences and Engineering Research Council of Canada (NSERC) and by Engage grant CGO 98069 from the NSERC. This support is gratefully acknowledged. We would also like to thank Louis Leclerc, Operations Manager, and Steve Thiboutot, Information Systems Manager, from Ameublement Tanguay, for their constant availability and comments, and for providing us with relevant data. Their support was critical to the success of this research and is greatly appreciated.

Results and views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect those of CIRRELT.

Les résultats et opinions contenus dans cette publication ne reflètent pas nécessairement la position du CIRRELT et n'engagent pas sa responsabilité.

* Corresponding author: Leandro.Coelho@cirrelt.ca

1 Introduction

In this paper we formally introduce, model and solve the vehicle routing problem with pauses (VRPP). This problem arises, e.g., in retail delivery, parcel and mail delivery, waste collection, and home blood sampling collection [27, 28, 22, 18], but we are not aware of any paper which has considered pauses explicitly. In the VRPP, drivers must make a pause during their shift to abide to law, union rules, and/or to company-specific regulations. The VRPP is a generalization of several classes of the well-known VRP [20]. Pauses are often seen in long-haul transportation, where drives are required to stop and rest after a given number of driving time to avoid excessively long working hours [12, 10, 13, 11, 14, 16, 15]. In last mile distribution, pauses are also present, for example, when drivers have time allotted for lunch breaks. However, although rest time is easily taken into account in long haul transportation (i.e., by adding the rest time to the trip duration), it is not the case for less than truckload routing as we will show later on.

Regarding industrial applications of the VRP, only a few works report results based on real cases and even fewer report industrial implementations. Among them, the food and soft-drink industries are the most prominent [29, 30, 17, 24, 6]. Other applications arise in different industries, including lubrication oil distribution [26, 31], waste collection [19, 4], industrial gases [8, 9], petroleum products [3, 5], and cash [32]. However, even if these articles deal with real applications, drivers pauses were not considered.

Our study is motivated by the request of an industrial partner facing this problem. Thus, we model, solve and evaluate drivers pauses as a real-world constraint for the first time. Despite a large body of research on VRPs, many industrial applications remain open and in some cases can only be handled with heuristics, due to the lack of a mathematical formulation capable of handling them. One of these constraints is related to drivers pauses. To the best of our knowledge, driver pauses have never been integrated into a mathematical formulation, and they have rarely been considered by existing heuristics. Notwithstanding, thousands of miles of drivers routes are often planned every day. In

practice, routes are planned disregarding the pause, which is later inserted into the final solution, yielding a poor approximation. This strategy can be applied for problems in which the pause separates two delivery periods, similar to a multi-period problem [1]. However, in a short planning horizon such as daily delivery in which our problem is defined, the pause does not lead to a multi-period problem. In fact, this strategy does not suit the retail industry where many customer deliveries should be done within some (tight) time windows and where service quality is highly important. In that context the timing of the pause becomes crucial.

The main contributions of this paper are threefold. First, we provide the first formulation, as a linear programming model for the VRPP. Second, we develop a fast and efficient heuristic capable of handling this practical constraint. Third, based on this heuristic we demonstrate that solving distribution problems with time windows without explicitly considering drivers pauses can lead to a number of infeasibilities in the final delivery schedules.

The remainder of this paper is organized as follows. In Section 2 we formally describe the VRPP. In Section 3 we introduce a linear programming formulation to model the VRPP. The description of our heuristic is presented in Section 4, which is followed by the computational experiments in Section 5. Conclusions are presented in Section 6.

2 Description of the problem

The VRPP is defined on a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where $\mathcal{V} = \{0, \dots, n\}$ is the vertex set and \mathcal{A} is the arc set. Vertex 0 corresponds to the depot, while the remaining vertices of $\mathcal{V}' = \{1, \dots, n\}$ represent the customers requesting a delivery. Each arc $(i, j) \in \mathcal{A}$ is associated with a travel time t_{ij} and a travel distance d_{ij} . Each customer i is associated with a service time S_i , a delivery weight w_i , and a delivery volume v_i . The service of a customer i must start within a time window $[a_i, b_i]$.

A heterogeneous fleet of K vehicles is located at the depot. Each vehicle k is associated with a volume capacity V_k and a weight capacity W_k . Vehicles routes must begin between $[a_s, b_s]$ and must end before a_e and, within these limits, each route must respect a maximum working time L_k . Vehicles are allowed to wait between two customers, and a lunch period lasting S_p units of time must be scheduled to start between $[a_p, b_p]$ in each route.

In this context, the goal is to minimize the total length of the routes in terms of traveling distance. Achieving reductions in traveling distance should lead to a reduction in the number of routes, which, in turn, can potentially decrease other important costs related to driver salaries and to the fixed and variable operating costs of the fleet.

In the next section we propose a mathematical formulation to solve the VRPP. To the best of our knowledge, this is the first time that such a problem is formulated.

3 Mathematical formulation

Since drivers' pauses can be taken at a customer site right after the service has been completed, or at a customer site just before starting the service, distances from the pauses' sites to the customers locations need to consider both cases. We provide in Sections 3.1 and 3.2 a detailed description of the transformations required to account for a pause in each route, and in Section 3.3 a mathematical formulation for the VRPP.

3.1 Extended distance matrix

Consider an instance with n customers, and its distance matrix d_{ij} defined over \mathcal{V}^2 . The distance matrix d_{ij} contains $n + 1$ rows and $n + 1$ columns. In order to account for pause and time windows, we will work with an extended distance matrix \tilde{d}_{ij} containing $2n + 2$ rows and columns. For notation purposes, we define a new set \mathcal{W} containing all nodes of $\mathcal{V} \cup \{n + 1, \dots, 2n + 2\}$. In the extended distance matrix \tilde{d}_{ij} , indices $n + 1$ to $2n + 1$ refer to pause nodes, one for each original node, and node $2n + 2$ indicates the depot.

$$d_{ij} = \begin{bmatrix} d_{00} & d_{01} & \cdots & d_{0n} \\ d_{10} & d_{11} & \cdots & d_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{n0} & d_{n1} & \cdots & d_{nn} \end{bmatrix}$$

$$d'_{ij} = \begin{bmatrix} \boxed{\begin{matrix} d_{00} & d_{01} & \cdots & d_{0n} \\ d_{10} & d_{11} & \cdots & d_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{n0} & d_{n1} & \cdots & d_{nn} \end{matrix}} & \boxed{\begin{matrix} d_{0,n+1} & d_{0,n+2} & \cdots & d_{0,2n+1} \\ d_{1,n+1} & d_{1,n+2} & \cdots & d_{1,2n+1} \\ \vdots & \vdots & \ddots & \vdots \\ d_{n,n+1} & d_{n,n+2} & \cdots & d_{n,2n+1} \end{matrix}} & \boxed{\begin{matrix} d_{0,2n+2} \\ d_{1,2n+2} \\ \vdots \\ d_{n,2n+2} \end{matrix}} \\ A_1 & A_2 & A_5 \\ \boxed{\begin{matrix} d_{n+1,0} & d_{n+1,1} & \cdots & d_{n+1,n} \\ d_{n+2,0} & d_{n+2,1} & \cdots & d_{n+2,n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{2n+1,0} & d_{2n+1,1} & \cdots & d_{2n+1,n} \end{matrix}} & \boxed{\begin{matrix} d_{n+1,n+1} & d_{n+1,n+2} & \cdots & d_{n+1,2n+1} \\ d_{n+2,n+1} & d_{n+2,n+2} & \cdots & d_{n+2,2n+1} \\ \vdots & \vdots & \ddots & \vdots \\ d_{2n+1,n+1} & d_{2n+1,n+2} & \cdots & d_{2n+1,2n+1} \end{matrix}} & \boxed{\begin{matrix} d_{n+1,2n+2} \\ d_{n+2,2n+2} \\ \vdots \\ d_{2n+1,2n+2} \end{matrix}} \\ A_3 & A_4 & A_6 \\ \boxed{\begin{matrix} d_{2n+2,0} & d_{2n+2,1} & \cdots & d_{2n+2,n} \\ d_{2n+2,n+1} & d_{2n+2,n+2} & \cdots & d_{2n+2,2n+1} \\ d_{2n+2,2n+2} \end{matrix}} & A_7 \end{bmatrix}$$

Obviously, the upper-left part of matrix d'_{ij} , identified as A_1 , corresponds to d_{ij} :

$$d'_{ij} = d_{ij} \quad i, j \in \mathcal{V}. \quad (1)$$

In order to model the problem, one needs to represent each location as two different nodes to adequately account for pauses. The reason is that there is no cost related to traveling to a pause location, but we must keep track of the last node visited before the pause in order to compute the appropriate distance to the next node after the pause. Thus, node $n + 1$ models a pause made immediately after leaving the depot (node 0); node $n + 2$ models a pause made immediately after visiting node 1, and so on. Since there is no cost for traveling to the pause, distances in the submatrix A_2 of d'_{ij} are modeled as follows:

$$d'_{ij} = \begin{cases} 0 & \text{if } j = i + n + 1 \\ \infty & \text{otherwise} \end{cases} \quad i \in \mathcal{V} \quad j \in \mathcal{W} \setminus (\mathcal{V} \cup \{2n + 2\}). \quad (2)$$

Then, one needs to model the distance from each pause node back to each customer $i \in \mathcal{V}$. If the vehicle visits the pause node associated with customer i , it cannot go back to i , thus the associated distance will be infinity. However, the vehicle can travel to any other customer j with the same original distance d_{ij} . Thus, the elements of d'_{ij} in the submatrix A_3 are modeled as:

$$d'_{ij} = \begin{cases} \infty & \text{if } j = i - n - 1 \\ d_{i-n-1,j} & \text{otherwise} \end{cases} \quad i \in \mathcal{W} \setminus (\mathcal{V} \cup \{2n + 2\}) \quad j \in \mathcal{V}. \quad (3)$$

Submatrix A_4 in d'_{ij} indicates an arc linking two pause nodes, which is forbidden. Thus, all its elements are equal to infinity:

$$d'_{ij} = \infty \quad i, j \in \mathcal{W} \setminus (\mathcal{V} \cup \{2n + 2\}). \quad (4)$$

Finally, one needs to model the last row and last column of the extended distance matrix, which refers to trips from and to the arrival depot node. Submatrix A_5 refers to trips from customers $i \in \mathcal{V}$ to the depot node $2n + 2$, which are all equal to d_{i0} :

$$d'_{i,2n+2} = d_{i0} \quad i \in \mathcal{V}. \quad (5)$$

Submatrix A_6 is related to trips from pause nodes to the arrival depot, which have the same distance as the trips from the original nodes associated with each pause node:

$$d'_{i+n+1,2n+2} = d_{i0} \quad i \in \mathcal{V}. \quad (6)$$

The last row of the extended distance matrix d'_{ij} refers to trips from the arrival depot

node to all other nodes. Since the arrival depot node is a sink, no vehicle is allowed to leave it, and all the elements in submatrix A_7 are equal to infinity:

$$d_{2n+2,j} = \infty \quad i \in \mathcal{W}. \quad (7)$$

Traveling times need to be updated in a similar fashion as the distance matrix.

3.2 Updating time windows and service duration parameters

The creation of the extended distance matrix leads to larger time window and service duration matrices as well. These updates must take into account according to the node they refer to (customers, pauses, or depot).

Using the nomenclature defined in the previous section, we now provide the time windows and service duration information for all nodes in \mathcal{W} . Nodes 0 to n refer to the original depot and customers of the problem, and thus their time windows and service duration remain unchanged.

All pause nodes are identified by indices $n + 1$ to $2n + 1$, and their time windows are all equal to $[a_p, b_p]$. The service duration when these nodes are visited is always equal to S_p .

The last node $2n + 2$ indicates the arrival depot, thus its service time is equal to zero. Each vehicle must return to the depot by a_e , so the time window of the returning depot is $[a_s, a_e]$.

3.3 Linear programming formulation

Using the extended distance matrix, we define an extended arc set \mathcal{Z} , with arcs $(i, j) \in \mathcal{Z}$, $i, j \in \mathcal{W}$. We formulate the VRPP using the following variables. Binary routing variables x_{ij}^k are equal to one if and only if arc $(i, j) \in \mathcal{Z}$, is used on the route of vehicle k . Binary variables y_i^k are equal to one if and only if node $i \in \mathcal{W}$ is visited by vehicle k . Continuous

variables s_i^k represent the starting time of the service for customer i by vehicle k . The problem is then formulated as follows:

$$\text{minimize } \sum_{(i,j) \in Z} \sum_{k \in \mathcal{K}} d_{ij} x_{ij}^k \quad (8)$$

subject to the following constraints:

$$\sum_{i \in \mathcal{V}'} v_i y_i^k \leq V_k y_0^k \quad k \in \mathcal{K} \quad (9)$$

$$\sum_{i \in \mathcal{V}'} w_i y_i^k \leq W_k y_0^k \quad k \in \mathcal{K} \quad (10)$$

$$\sum_{k \in \mathcal{K}} y_i^k = 1 \quad i \in \mathcal{V}' \quad (11)$$

$$\sum_{i \in \mathcal{W} \setminus (\mathcal{V} \cup \{2n+2\})} y_i^k \leq 1 \quad k \in \mathcal{K} \quad (12)$$

$$\sum_{i \in \mathcal{W} \setminus (\mathcal{V} \cup \{2n+2\})} y_i^k = y_0^k \quad k \in \mathcal{K} \quad (13)$$

$$n y_0^k \geq \sum_{i \in \mathcal{V}'} y_i^k \quad k \in \mathcal{K} \quad (14)$$

$$\sum_{i \in \mathcal{W} \setminus \{0, 2n+2\}} x_{0i}^k = y_0^k \quad k \in \mathcal{K} \quad (15)$$

$$y_{2n+2}^k = y_0^k \quad k \in \mathcal{K} \quad (16)$$

$$\sum_{i \in \mathcal{W} \setminus \{0, 2n+2\}} x_{i, 2n+2}^k = y_{2n+2}^k \quad k \in \mathcal{K} \quad (17)$$

$$\sum_{j \in \mathcal{W}} x_{ij}^k + \sum_{j \in \mathcal{V}} x_{ji}^k = 2y_i^k \quad i \in \mathcal{W} \setminus \{0, 2n+2\} \quad k \in \mathcal{K} \quad (18)$$

$$\sum_{i \in \mathcal{W}} x_{ij}^k \leq 1 \quad i \in \mathcal{W} \quad k \in \mathcal{K} \quad (19)$$

$$\sum_{i \in \mathcal{W}} x_{ji}^k \leq 1 \quad i \in \mathcal{W} \quad k \in \mathcal{K} \quad (20)$$

$$s_i^k + S_i + l_{ij} - M(1 - x_{ij}^k) \leq s_j^k \quad i, j \in \mathcal{W} \quad k \in \mathcal{K} \quad (21)$$

$$s_i^k \geq a_i \quad i \in \mathcal{W} \quad k \in \mathcal{K} \quad (22)$$

$$s_i^k \leq b_i \quad i \in \mathcal{W} \quad k \in \mathcal{K} \quad (23)$$

$$s_{2n+2}^k - s_0^k \leq L_k \quad k \in \mathcal{K} \quad (24)$$

$$x_{ij}^k, y_i^k \in \{0, 1\} \quad (i, j) \in \mathcal{Z} \quad i \in \mathcal{W} \quad k \in \mathcal{K}. \quad (25)$$

The objective function (8) aims at minimizing the total vehicle routing costs. Constraints (9) and (10) impose vehicle capacities with respect to the total volume and weight, respectively. Constraints (11) impose that all customers must be visited by exactly one vehicle. Constraints (12) impose that at most one pause node is visited by each vehicle, while constraints (13) require that a pause node is visited if the vehicle is used. Constraints (14) impose that the vehicle visits the depot if any customer is assigned to it, and constraints (15) ensure that the vehicle leaves the depot. Constraints (16) ensure that if a vehicle leaves the depot node 0 it should return to the depot node $2n+2$. Constraints (17) ensure that if the vehicle should return to the depot, one arc towards it is used. Constraints (18) are degree constraints, and constraints (19) and (20) ensure that there are at most one incoming and one outgoing arc for each node. Time windows and subtour elimination are imposed by means of constraints (21). Bounds on the time windows for the beginning of the service on every node are imposed through constraints (22) and (23). Shift duration constraints are imposed through constraints (24). Integrality conditions are imposed by constraints (25).

4 Heuristic routing algorithm

We will show later that the formulation presented in Section 3.3 is too difficult to be solved even for small sized instances of VRPP. Thus, we developed an insertion based heuristic, which will help us to demonstrate the impact of managing driver pauses in larger instances inspired from real data.

The routing algorithm contains four main procedures, which are repeated a predetermined number of times (20 in our implementation). We describe the initialization phase in

Section 4.1, and the procedure to route customers in Section 4.2. A local improvement procedure is described in Section 4.3, and a route compression phase is presented in Section 4.4.

An important procedure in our algorithm is related to the way we open new routes. We describe in Section 4.5 the procedure which opens one or several new routes as required.

In our algorithm, pauses are handled by adding a dummy customer with a time window $[a_p, b_p]$ and a service time S_p to each route. The distances between dummy customers and other customers are set to zero.

It is worth mentioning that because of the practical nature of the problem, and due to the limited number of vehicles and to tight customer time windows, it may not be possible to include all customers in a solution. These unserved customers are referred to as *exceptions*. In practice, the list of exceptions is handled by the customer service department who will contact them in order to schedule another delivery date. Such a set of customers not to be visited has already been treated in the literature [7, 2].

4.1 Initialization

The first phase of the algorithm creates and initializes the different data sets that will be used during the routing process. Let \mathcal{R} be the set of active routes and \mathcal{E} be the set of exceptions, both initially empty. Let \mathcal{C} be a set containing all customers. The following procedure is executed in order to open one or several routes and to sequentially assign customers in \mathcal{C} until all the customers have been assigned to routes or to the set \mathcal{E} .

4.2 Routing customers

This step in our algorithm plans visits to customers from \mathcal{C} into the existing routes. At the end of this step, all customers will be assigned to a route or to the exception set \mathcal{E} . For each customer, we use the following procedure to find the best route and to place the

customer in the best position.

We try to insert each customer in each possible position of all the existing routes. The insertion between two customers is feasible only when it respects capacity constraints, all time windows, total route duration, and end of route time. When all the routes and all the insertion positions have been tested, the procedure returns for each customer the best insertion position and cost (see line 7 of Algorithm 1). The next customer to be inserted is obtained by dividing the insertion cost by a random value taken from the interval $[\alpha, 1]$ (see lines 8–11 of Algorithm 1), where $\alpha = 1 - \frac{\text{iteration}-1}{\text{max iterations}}$, and selecting the customer with the smallest value. This means that the first iteration is fully deterministic, while the remaining successive iterations are each more randomized than the previous one.

If an insertion is possible, we remove the selected customer from \mathcal{C} , update \mathcal{R} , and repeat the procedure. Otherwise, we try to open a new route using the procedure described in Section 4.5. If no more vehicles are available, the remaining customers are added to the exception set \mathcal{E} . We provide a sketch of this procedure in Algorithm 1.

4.3 Local improvement

The local improvement phase consists of two neighborhood search heuristics which are applied to all routes in \mathcal{R} in order to decrease their total length. We first apply an intra-route search, which is described in Section 4.3.1, followed by an inter-routes search, described in Section 4.3.2.

Two feasibility checks are performed in order to accept a move. First, we evaluate the vehicle capacity, both in terms of volume and weight, as well as the total route duration. The second feasibility check concerns the time windows of all the customers in the route.

Algorithm 1 Routing customers (Section 4.2)

```

1: best_client = 0
2: best_ratio =  $\infty$ 
3: if  $\mathcal{C} = \emptyset$  then
4:   Go to 26.
5: end if
6: for  $i \in \mathcal{C}$  do
7:    $c_i \leftarrow$  best feasible insertion for customer  $i$ 
8:    $c'_i \leftarrow \frac{c_i}{rand[\alpha, 1]}$ 
9:   if  $c'_i < \text{best\_ratio}$  then
10:    best_ratio  $\leftarrow c'_i$ 
11:    best_client  $\leftarrow i$ 
12:   end if
13: end for
14: if best_client  $\neq 0$  then
15:   Insert best_client in its best position
16:    $\mathcal{C} \leftarrow \mathcal{C} \setminus \{\text{best\_client}\}$ 
17:   Go to 1.
18: else
19:   if there are vehicles available then
20:    Call open new route(s) procedure (see Section 4.5)
21:    Go to 1.
22:   else
23:    Insert all customers from  $\mathcal{C}$  in  $\mathcal{E}$ 
24:   end if
25: end if
26: Return routes and  $\mathcal{E}$ .

```

4.3.1 Intra-route improvement

The intra-route improvement considers one route at a time. It applies the 3-opt algorithm of Lin [21] in which we also check for time window feasibility, until no more improvements can be obtained.

4.3.2 Inter-routes improvement

This procedure considers all routes in $\mathcal{R} = \{r_1, \dots, r_K\}$ and tries to improve each pair of routes r_i ($i = 1, \dots, K-1$) and r_j ($j = i+1, \dots, K$) by exchanging customers between them. As proposed by Renaud and Boctor [25] for the fleet size and mixed vehicle routing problem, 11 exchanges are evaluated. These exchanges are detailed in the Appendix and are applied to all possible chains of four consecutive vertices between routes r_i and r_j . The depot and the dummy nodes representing the drivers' pauses are never exchanged. The first improving move is applied and the procedure continues as long as improvements are obtained. Once all the possible pairs of routes have been considered, the procedure restarts by evaluating only the pairs of routes for which at least one of the routes have been modified during the previous iteration. This procedure can be viewed as a restriction of the 2-interchange procedure of Osman [23], however we concentrate on more promising moves yielding a much faster heuristic.

4.4 Route compression

The last phase of our algorithm is a route compression procedure which aims at reducing the total time duration of a route. All insertion procedures described so far minimize the traveling distance and try to insert customers as soon as possible along the route, so that it is possible that some routes include unnecessary waiting times.

The route compression procedure tries to postpone visits as much as possible. It starts by setting the visit to the last customer as soon as possible in order to finish the workday

earlier. Then, moving from the back of the route towards its beginning, each visit is rescheduled to as late as possible, still respecting the time window constraints.

In addition to the potential reduction in the total route duration, this compression procedure also has two positive side effects. First, the visit to the last customer is performed as soon as possible, which is usually appreciated by the customers and drivers. Second, the first visit is performed as late as possible, which besides being appreciated by the customers also helps drivers avoid the early morning traffic.

4.5 Opening new routes

One important feature of our algorithm is that, unlike other classic approaches which open one new route at a time, it can open up to π new routes simultaneously whenever the existing routes cannot accommodate the remaining customers in \mathcal{C} . The motivation for opening several routes instead of only one is to avoid having very busy routes which are concentrated in a specific area, which leads to very high costs when customers away from this cluster need to be visited. Thus, if two or more routes are open, the procedure tries to initialize them in different directions aiming at a better geographic coverage.

The number of routes to open depends on the number of the remaining customers to be visited, and on their needs. In particular, the procedure computes an estimation on the number of routes required to fulfill the deliveries to the remaining customers. This estimation is based on three ratios regarding the capacity of the vehicles in terms of both volume and weight, and the traveling and service times. The first two ratios are computed as the total required volume (weight) of the deliveries divided by the average capacity of the vehicles. The third one is based on the maximum shift duration, and is equal to the ratio between the estimation of the total time required to visit all the remaining customers and the maximum shift duration. The total time required to visit all customers is estimated as the average traveling time between each unvisited customer i and all the nodes in $\mathcal{C} \cup \{0\}$ plus their required service times. The maximum of these three ratios

provides the number of routes to be opened by this procedure, bounded by the number of available vehicles and by a parameter indicating the maximum number of routes to be opened per iteration. New routes are opened using the vehicles with larger capacity.

Each of the new routes is initialized with a pause and by adding a customer to it in the following manner. If only one route is created, we include in this route the customer in \mathcal{C} the farthest customer from the depot. If two routes are opened, two customers i and j in \mathcal{C} are selected in such a way that the total distance $d_{0i} + d_{ij} + d_{j0}$ is maximized. If three routes are created, three customers i , j and k are selected such that $d_{0i} + d_{0j} + d_{0k} + d_{ij} + d_{jk} + d_{ki}$ is maximized. We have limited π to three as per our preliminary numerical experiments. This way, new routes will tend to cover opposite regions.

Each of the new routes are added to \mathcal{R} , and the customers that have been assigned to them are removed from \mathcal{C} .

5 Computational experiments

In this section we describe the results of the computational experiments carried out to show the relevance of adequately planning drivers' pauses, and to evaluate the performance of the proposed mathematical formulation and of the heuristic. First, we show that the VRPP model cannot be solved even for very small instances. Second, we demonstrate that driver pauses should be considered in the optimization phase, and that neglecting this step, for example, by creating vehicle routes without pauses which are reinserted later, can lead to very poor solutions.

Our heuristic was implemented using VB.net (.net framework 4.0). Numerical experiments were carried out on a desktop equipped with an Intel Core i7-3612QM @ at 2.1 GHz and with 8 GB of RAM running MS Windows x64. The linear programming formulation was implemented in C++ using CPLEX Concert Technology and the experiments were carried out on a desktop equipped with an Intel i7 running at 3.66 GHz and with 8

GB of RAM, under a Linux operating system.

We start by comparing the performance of the VRPP formulation proposed in Section 3 to the one of the heuristic presented in Section 4. We have generated 15 small instances based on a large real-life one obtained from an industrial partner. These instances contain from five to 20 delivery requests, and were solved by both methods. Table 1 presents average results, and shows that the branch-and-bound algorithm quickly becomes inefficient when the size of the instance increases. In particular, we observe that after one hour of computing time, the branch-and-bound algorithm is not able to find optimal solutions, nor to provide reasonable gaps for instances containing 20 requests. Given that the real instances provided by our partner contain at least one hundred requests and can go up to 400 requests, it becomes clear that the proposed formulation is not suitable for real applications. Table 1 also shows that the computing time required by our heuristic remains extremely low when the size of the instances increases. Moreover, on these small instances for which the exact algorithm was able to obtain optimal solutions, our heuristic performed very well, being able to yield optimal solutions when these are known. This enables us to further evaluate the impact of the pauses in distribution problems.

Table 1: Comparison of the performance of the branch-and-bound algorithm for the VRPP and of the heuristic (average over five instances inspired by real data)

# requests	B&B algorithm			Heuristic	
	Distance	Gap (%)	Time (s)	Distance	Time (s)
5	182.6	0.00	3.2	182.6	0.178
10	282.4	0.00	737.6	282.4	0.497
20	733.6	51.58	3600.0	709.8	1.365

We have shown that explicitly considering drivers pauses greatly increases the size of the model, so that only very small instances can be solved to optimality. The most intuitive workaround to adapt classical algorithms to obtain solutions for the VRPP is to solve the problem in two phases as follows. First, one solves the problem without considering

the pauses by using well known VRP algorithms. Then, one reintroduces the pause in the best possible position and adjust the delivery schedules accordingly. We now aim at showing how this two-phase approach for handling the pauses can lead to major solution infeasibilities. Moreover, we empirically demonstrate that the number of infeasibilities rapidly increases when the percentage of deliveries with time windows increases, and also when time windows are tighter. A real instance with 209 deliveries provided by our industrial partner is used to support our experiments plan.

First, we removed all time windows from the instance. Then we randomly generated time windows for 10%, 20%, 30%, 40% and 50% of the deliveries. These time windows were randomly generated between 8:30 and 15:00. It is worth to mention that, since the drivers' pause must be taken between 11:30 and 13:00, morning time windows are not impacted by the pause, but they still help shape the final solution. We have repeated this experiment twice, first considering 3-hours time windows, and then tighter 2-hours time windows. We solved these instances by the two-phase approach just described. To this end, we first set the pause time duration to zero and reduced the length of the day accordingly, i.e., drivers should return to the depot one hour earlier. The problem was then solved without pauses, using the heuristic described in Section 4. In the resulting solution, a pause with time zero was included in each route. In the second phase, the pause length was set back to 60 minutes and the driver schedule after the pause was updated. We also tried to relocate the pause to all possible positions within its time window to reduce infeasibilities as much as possible. If reinserting the pause leads to an infeasible solution, the position leading to the fewer violated time windows was selected. If many feasible positions were available, the one which yielded the shortest route was selected. A summary of these experiments is provided in Tables 2 and 3, where the column *% of TW* shows the percentages of deliveries having a time window, *Km* and *# of routes* shows the number of kilometers and the number of routes in the solution, respectively, the column *# of infeasible routes* provides the number of infeasible routes in the solution, and *# of infeasible customers* shows the total number of customers that are visited outside of their

time windows.

Table 2: Impact of driver pauses with three-hour time windows

% of TW	Km	# of routes	# of infeasible routes	# of infeasible customers
0	2913	22	—	—
10	3290	23	3	6
20	3325	23	3	4
30	3059	23	7	12
40	3044	24	6	8
50	3089	24	10	18

Results reported in Table 2 clearly show the impact of neglecting driver pauses in the solution process when three-hour time windows are imposed. Even with only 30% of deliveries having time windows, seven routes out of 23 (30%) are infeasible, with 12 late deliveries (5.7%). When 50% of the requests contain time windows, 41% of the routes become infeasible and 18 deliveries are late (8.6%). If we consider that, on average, half of 209 deliveries are performed in the afternoon, this means in fact that 18 deliveries out of 104.5 are late, which rise the real percentage of late deliveries to 17.2%.

Table 3 presents results with tighter, two-hour, time windows. As expected, the results quickly deteriorate. For 30% and 50% of deliveries with time windows, the number of infeasible deliveries nearly doubled to 20 and 38 customers when comparing instances with 2- and 3-hour time windows. These results clearly demonstrate that driver pauses should be directly included in the resolution process and that not doing so yields a bad approximation which leads to numerous infeasibilities.

Table 3: Impact of driver pauses with two-hour time windows

% of TW	Km	# of routes	# of infeasible routes	# of infeasible customers
0	2913	22	—	—
10	2888	23	3	6
20	3031	23	7	8
30	2862	23	11	20
40	3163	23	16	32
50	3235	24	16	38

6 Conclusions

We have formally introduced the vehicle routing problem with pauses (VRPP). We have proposed the first linear programming formulation for this problem and showed that a branch-and-bound algorithm applied to it is impractical for real-life applications. We have also developed an efficient insertion heuristic which is capable of solving instances containing several hundred customers with time windows, several vehicles, and places the drivers pauses in the required time window. The results of our computational experiments confirm that one must solve the problem considering the drivers pauses in the optimization algorithm, at the risk of not being able to obtain feasible solutions for many situations.

Appendices

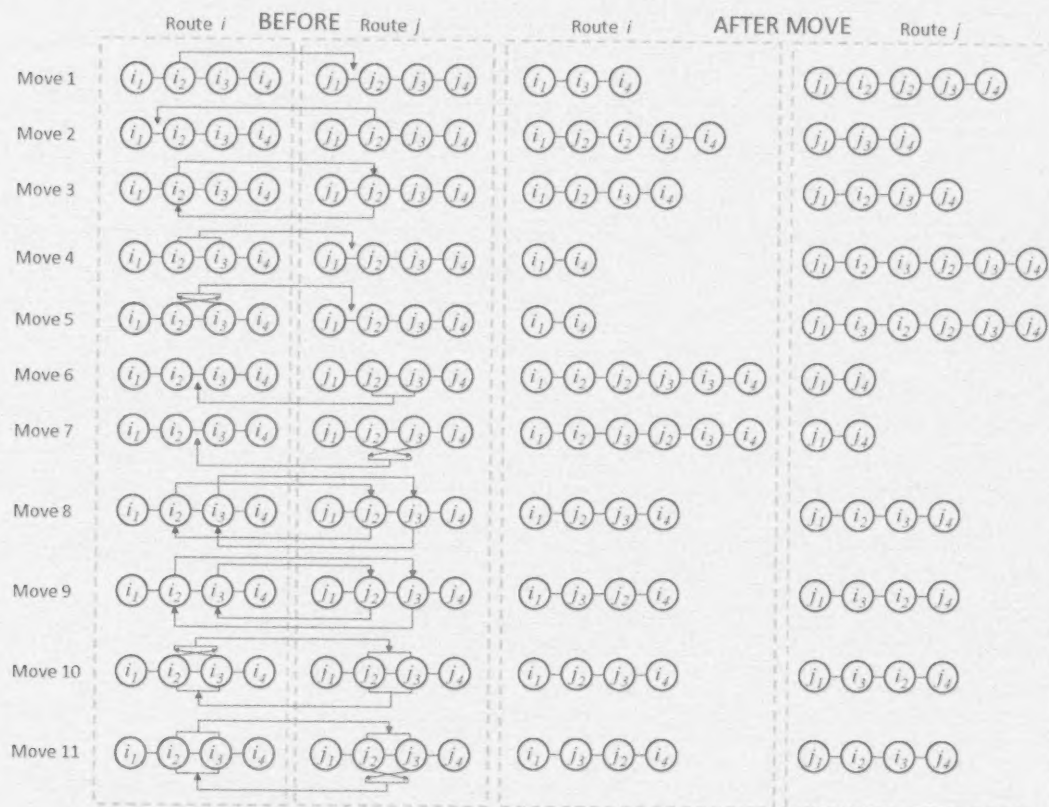
Intra-route movements

Let (i_1, i_2, i_3, i_4) and (j_1, j_2, j_3, j_4) be two chains of four vertices from routes r_i and $r_j \in \mathcal{R}$. The following 11 moves are considered:

1. Place i_2 between j_1 and j_2 .

2. Place j_2 between i_1 and i_2 .
3. Swap i_2 and j_2 .
4. Move i_2 and i_3 to r_j , inserting (i_2, i_3) between j_1 and j_2 .
5. Move i_2 and i_3 to r_j , inserting (i_3, i_2) between j_1 and j_2 .
6. Move j_2 and j_3 to r_i , inserting (j_2, j_3) between i_1 and i_2 .
7. Move j_2 and j_3 to r_i , inserting (j_3, j_2) between i_1 and i_2 .
8. Swap i_2 and j_2 , swap i_3 and j_3 .
9. Swap i_2 and j_3 , swap i_3 and j_2 .
10. Replace i_2 and i_3 by j_2 and j_3 , respectively, and replace j_2 and j_3 by i_3 and i_2 , respectively.
11. Replace i_2 and i_3 by j_3 and j_2 , respectively, and replace j_2 and j_3 by i_2 and i_3 , respectively.

Moves 1 to 3 correspond to the 1-interchange procedure of Osman [23], and moves 4 to 11 represent a subset of the 2-interchange exchanges tested by Osman [23].

Figure 1: Representation of the 11 moves evaluated during the intra-route improvement phase

References

- [1] C. Archetti and M. Savelsbergh. The trip scheduling problem. *Transportation Science*, 43(4):417–431, 2009.
- [2] C. Archetti, A. Hertz, and M. G. Speranza. Metaheuristics for the team orienteering problem. *Journal of Heuristics*, 13(1):49–76, 2007.
- [3] P. Avella, M. Boccia, and A. Sforza. Solving a fuel delivery problem by heuristic and exact approaches. *European Journal of Operational Research*, 152(1):170–179, 2004.

- [4] A. M. Benjamin and J. E. Beasley. Metaheuristics for the waste collection vehicle routing problem with time windows, driver rest period and multiple disposal facilities. *Computers & Operations Research*, 37(12):2270–2280, 2010.
- [5] F. F. Boctor, J. Renaud, and F. Cornillier. Trip packing in petrol stations replenishment. *Omega*, 39(1):86–98, 2011.
- [6] M. Caramia and F. Guerriero. A milk collection problem with incompatibility constraints. *Interfaces*, 40(2):130–143, 2010.
- [7] I.-M. Chao, B. L. Golden, and E. A. Wasil. The team orienteering problem. *European Journal of Operational Research*, 88(3):464–474, 1996.
- [8] J. M. Day, P. D. Wright, T. Schoenherr, M. Venkataramanan, and K. Gaudette. Improving routing and scheduling decisions at a distributor of industrial gasses. *Omega*, 37(1):227–237, 2009.
- [9] K. C. Furman, J.-H. Song, G. R. Kocis, M. K. McDonald, and P. H. Warrick. Feedstock routing in the ExxonMobil downstream sector. *Interfaces*, 41(2):149–163, 2011.
- [10] A. Goel. Vehicle scheduling and routing with drivers' working hours. *Transportation Science*, 43(1):17–26, 2009.
- [11] A. Goel. Truck driver scheduling in the european union. *Transportation Science*, 44(4):429–441, 2010.
- [12] A. Goel and V. Gruhn. Drivers' working hours in vehicle routing and scheduling. In *Intelligent Transportation Systems Conference, 2006*, pages 1280–1285, 2006.
- [13] A. Goel and L. Kok. Truck driver scheduling in the United States. *Transportation Science*, 46(3):317–326, 2010.
- [14] A. Goel and L.-M. Rousseau. Truck driver scheduling in Canada. *Journal of Scheduling*, 15(6):783–799, 2012.

- [15] A. Goel and T. Vidal. Hours of service regulations in road freight transport: An optimization-based international assessment. *Transportation Science*, 2013. doi: 10.1287/trsc.2013.0477.
- [16] A. Goel, C. Archetti, and M. Savelsbergh. Truck driver scheduling in Australia. *Computers & Operations Research*, 39(5):1122–1132, 2012.
- [17] B. L. Golden, A. A. Assad, and E. A. Wasil. Routing vehicles in the real world: applications in the solid waste, beverage, food, dairy, and newspaper industries. In P. Toth and D. Vigo, editors, *The Vehicle Routing Problem*, Monographs on Discrete Mathematics and Applications, pages 245–286. SIAM, Philadelphia, 2002.
- [18] J. Gromicho, J. J. van Hoorn, A. L. Kok, and J. M. J. Schutten. Restricted dynamic programming: a flexible framework for solving realistic VRPs. *Computers & Operations Research*, 39(5):902–909, 2012.
- [19] B.-I. Kim, S. Kim, and S. Sahoo. Waste collection vehicle routing problem with time windows. *Computers & Operations Research*, 33(12):3624–3642, 2006.
- [20] G. Laporte. Fifty years of vehicle routing. *Transportation Science*, 43(4):408–416, 2009.
- [21] S. Lin. Computer solutions of the traveling salesman problem. *Bell System Technical Journal*, 44(10):2245–2269, 1965.
- [22] B. M. Ombuki-Berman, A. Runka, and F. T. Hanshar. Waste collection vehicle routing problem with time windows using multi-objective genetic algorithms. In *Proceedings of the Third IASTED International Conference on Computational Intelligence*, pages 91–97, 2007.
- [23] I. H. Osman. Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research*, 41(4):421–451, 1993.

- [24] J. Privé, J. Renaud, F. F. Boctor, and G. Laporte. Solving a vehicle-routing problem arising in soft-drink distribution. *Journal of the Operational Research Society*, 57(9): 1045–1052, 2005.
- [25] J. Renaud and F. F. Boctor. A sweep-based algorithm for the fleet size and mix vehicle routing problem. *European Journal of Operational Research*, 140(3):618–628, 2002.
- [26] P. P. Repoussis, D. C. Paraskevopoulos, G. Zobolas, C. D. Tarantilis, and G. Ioannou. A web-based decision support system for waste lube oils collection and recycling. *European Journal of Operational Research*, 195(3):676–700, 2009.
- [27] Y. Rochat and F. Semet. A tabu search approach for delivering pet food and flour in Switzerland. *Journal of the Operational Research Society*, pages 1233–1246, 1994.
- [28] Y. Rochat and É. D. Taillard. Probabilistic diversification and intensification in local search for vehicle routing. *Journal of heuristics*, 1(1):147–167, 1995.
- [29] C. D. Tarantilis and C. T. Kiranoudis. A meta-heuristic algorithm for the efficient distribution of perishable foods. *Journal of Food Engineering*, 50(1):1–9, 2001.
- [30] C. D. Tarantilis and C. T. Kiranoudis. Distribution of fresh meat. *Journal of Food Engineering*, 51(1):85–91, 2002.
- [31] M. F. Uzar and B. Çatay. Distribution planning of bulk lubricants at BP Turkey. *Omega*, 40(6):870–881, 2012.
- [32] R.G van Anholt, L. C. Coelho, G. Laporte, and I. F. A. Vis. An inventory-routing problem with pickups and deliveries arising in the replenishment of automated teller machines. Technical Report CIRRELT-2013-71, Montreal, Canada, 2013.